arXiv:1602.07586v2 [math.ST] 15 May 2016

A PARSIMONIOUS THEORY OF EVIDENCE-BASED CHOICE

FATEMEH BORHANI AND EDWARD J. GREEN
PENNSYLVANIA STATE UNIVERSITY

Abstract. Primitive entities of the theory to be presented here are a body of evidence available to an agent (called an evidential state) and an alternative in a set, from which the agent might choose. Assumptions are stated regarding the space of possible evidential states. Under those assumptions, while the space of evidential states is not necessarily a Boolean algebra, it can be embedded in a structure-preserving way into a canonical sigma-field of events. A plan is a mapping from evidential states to choice alternatives. A consistency condition for plans, reminiscent of Savage's sure-thing principle, is formulated. The condition is neither necessary nor sufficient for a plan to be rationalized by subjective-utility maximization with respect to a probability measure on the canonical sigma-field. A structure of evidential states may contain, or coincide with, a substructure that models a process of experimental learning. A plan specified on such a substructure satisfies the consistency condition if, and only if, it can be rationalized by maximization of subjective conditional expected utility.

1. Introduction

The research to be reported here provides an alternate perspective, from that provided by representation theorems such as those of de Finetti [1937] and Savage [1972], on the choices of an agent who may, or may not, maximize subjective expected utility (SEU). A model is constructed in which choice alternatives are primitive entities (instead of being "Savage acts" that are constructed from "states of the world" and "consequences"), and in which the evidence available to a decision maker is also a primitive entity that might be regarded, for example, as being a concrete record of experimental observations rather than as being a subset of all possible states of the world. A measurable space of states of the world is constructed as (in theorem 6) a derived entity of the theory, rather than being a primitive entity of the model. Choice, rather than preference, is modeled. Under

Date: 2016.05.06

.

The authors gratefully acknowledge the Human Capital Foundation, (`http://www.hcfoundation.ru/en/`) and particularly Andrey P. Vavilov, for research support through the Center for the Study of Auctions, Procurements, and Competition Policy (CAPCP, `http://capcp.psu.edu/`) at the Pennsylvania State University.

some assumptions, there is a revealed-preference characterization involving state-dependent utility of choice alternatives. Consequences play no role whatsoever in the theory.

This research carries further what we understand to be the fundamental principle that motivates recent work by Karni and his associates (for example, in Karni [2006, 2015]): that an entity that is best regarded as being in the mind of the decision maker should not be a primitive entity of a theory that is intended to explain an agent's observable choices in terms of the agent's observable situation. Karni applies this principle specifically to states of the world, provides arguments in [2006] for that application of it, and cites further discussion in the same vein by Gilboa and Schmeidler [2001]. We regard the principle as applying equally to consequences. In fact, both Savage and Karni would agree. Savage [1972, p. 84] wrote that "what are often thought of as consequences (that is, sure experiences of the deciding person)... typically are in reality highly uncertain... [and] can perhaps never be well approximated." Karni [2006, p. 326] writes that "Often the distinction between states and consequences is blurred[.]"

This perspective contrasts in numerous respects with the versions of Bayesian decision theory developed by de Finetti [1937] and Savage [1972], whose frameworks have largely been adopted in the most familiar generalizations of the theory of choice under uncertainty.[1] De Finetti's and Savage's contributions might be called *representation theories* of SEU maximization. The characterization to be developed here, under some specific assumptions regarding the structure of evidence, might be called a *revealed-preference theory* of SEU maximization. In fact, the concepts of utility representability of preferences and rationalizability of agents' plans that will be defined in sections 4–6 are essentially a specialization of the abstract revealed-preference framework developed by Richter [1966].

Revealed-preference theory grew out of the economic theory of the budget-constrained consumer, and there is a research literature on revealed preference for budget-constrained expected-utility maximization.[2] In contrast, rather than conceiving a situation to be a budget set from which the agent is permitted to choose a consumption bundle, our theory conceives the agent's situation in virtually the same terms as Savage's theory does.

Consider the following quotation from Savage.

> Inference means for [Bayesians] the changes in opinion induced by *evidence* on the application of Bayes' theorem. For us, a problem in analysis... can... be separated into two phases. First, compute the new [probability] distribution induced by the original distribution *and the data.* Second, use

[1] Savage [1961, section 6] succinctly describes the work of additional main contributors to the theory of subjective (or personal) probability. The theory to be presented in sections 4–7 builds on, and makes a correction to, Green and Park [1996].

[2] Cf. Polisson et al. [2015, `http://www.economics.ox.ac.uk/materials/papers/13804/paper740.pdf`] and the articles cited there.

> the new distribution and the economic facts about terminal actions to pick one of the terminal actions that now has the highest expected utility.[3]

In order to understand the import of this passage, let us shift focus for a moment from the agent to an investigator who studies the agent. How would Savage have answered, if he had been asked what sort of data set an investigator might use to determine whether or not an agent were behaving consistently with Bayesian rationality (that is, behaving consistently with the assumptions of having beliefs conforming to a prior subjective-probability measure, applying Bayes' theorem, and taking an action—or, in general, making a choice among alternatives—to maximize conditional expected utility)? We read Savage as meaning 'economic facts' in a broader sense than as a specific reference to budget-constrained choice, so the agent's terminal actions could be actions of any sort. But Savage is clear that the agent's situation (as it pertains to inference) is simply the body of evidence/data of which the agent is currently in possession.

That is, Savage might have answered that the investigator's data set is a list of observations, each of which specifies a body of evidence that the agent possesses at some point together with the action that the agent takes at that point. Moreover, since Savage calls the action 'terminal', we understand that, in the specific passage just quoted, he does not envision the agent as being in an interim stage of performing a sequence of actions chosen by dynamic programming. Rather, akin to what he had previously written about 'small worlds', Savage envisions that the evidential situation and chosen action or alternative delimit a compartment of the agent's life that can satisfactorily be considered in isolation.[4]

Broadly speaking, a revealed-preference theory has two goals that distinguish it from a representation theory.

(1) To describe an agent's behavior in terms that are scrupulously atheoretical, corresponding as closely as possible to how that behavior would be described in an actual data set. Such a description takes the form of a function (or a multi-valued correspondence), defined on a set of situations that may face the agent, that specifies what the agent would do, or would choose, in each of those situations.

(2) To formulate minimal conditions regarding the agent's behavior that imply consistency with a specified set of assumptions about (that is, with a theory of) rational choice; and, moreover, to investigate what further conditions would be required to ensure consistency if more demanding assumptions are added to the theory.

The relationship between the theory to be presented here and the theories of De Finetti and Savage, is a case in point. Regarding (1), the formalization of choice alternatives, or acts, exemplifies the difference in perspective

[3]Savage [1961, p. 178, emphasis added]

[4]Cf. Savage [1972, section 5.5].

between them. For example, if the weather forecast specified rain, then an agent might either stay at home, carry an umbrella while going outside, or go outside without an umbrella. Choice alternatives will be primitive entities in the present theory, so the intuitive descriptions just given would be used to identify or name them (e.g., $h =$ to stay at home, $u =$ to carry an umbrella, $r =$ to take the risk of rain) but would play no further role in a formal analysis. In Bayesian representation theory, an act is a derived, theory-laden entity. It is defined as a mapping from a Boolean algebra (or, in some versions of the theory, a σ-field) of events—each represented, in turn, as a set of unobserved states of the world—to consequences, which are primitive entities.[5] That is, a commitment to the theoretical entities of Bayesian decision theory is required before an act can even be specified at the first, descriptive, stage of theorizing. This is the sort of theoretical pre-commitment that it is a goal of revealed-preference theory to avoid. The results to be proved below also differ in spirit from representation theory with respect to goal (2). Theorem 24 establishes the existence, but not the uniqueness, of an expected-utility rationalization of behavior that meets prescribed conditions.[6] Other analysis to be presented here, especially example 22, concerns what features a pattern of behavior must possess in order for an expected-utility representation to satisfy conditions that are typically invoked in applied Bayesian analysis, but that are not logically implied by the criterion of Bayesian rationality. That is, the analysis to be presented here is weaker in one respect, but is more extensive in another respect, than the representation-theoretic analysis is.

Notation. References to objects indicate their level in the set theoretic hiearachy. Level-0 objects are numbers (i, j, k, m, n), evidential states (r and q–z), sample points (ϕ, χ, ω), and choice alternatives (a, b, c).

Upper-case, math-italic letters (for example, B) denote level-1 objects (sets or ordered pairs of level-0 objects). The sets of numbers, evidential states, sample points, and choice alternatives are named $\mathbb{N}$, $\mathbb{X}$, Ω, and A, respectively. Υ and Φ denote sets of sample points in a measurable space that is not the same as Ω. Letter ω will always refer to an element of Ω, while ϕ, and χ will refer to elements of Φ.

$\mathbb{N} = \{0, 1, 2, \dots\}$ and $\mathbb{N}_+ = \{1, 2, 3, \dots\}$.

[5] In the example, staying at home would be modeled as the act that never leads to getting wet or to meeting a celebrity in the street, while going out without an umbrella would be modeled as an act that might have any combination of those outcomes as a consequence in some state of the world.

[6] Representation theorems for state-contingent expected utility maximization, such as Karni et al. [1983], ensure uniqueness in the weak sense that the measures that can be derived from integration of the utility function with respect to the subjective probability measure form a two-dimension subset (corresponding to affine transformations of the utility scale) of the infinite-dimensional space of finite measures. Not even that sort of weak uniqueness will be derivable here.

Upper-case calligraphic letters (for example, $\mathcal{B}$) denote level-2 objects (sets of sets).

The letters σ and Σ, as well as the usual symbols (for example, $\leq$) denote order relations.

Functions from their domains to level-0 objects are denoted by lower-case math-italic and Greek letters (for example, f, ρ). Operators from their domains to higher-level objects are denoted by upper- and lower-case Fraktur letters. For example, $\mathfrak{F}(\mathcal{T})$ denotes the field of sets generated by the sets in topology $\mathcal{T}$ that are both open and closed, of subsets of Ω, and $\mathfrak{B}_T$ denotes the set of branches of a graph-theoretic tree, $(T, \leq)$. The mapping that assigns an event (that is, a set of sample points) to each evidential state is denoted by $\mathfrak{e}$ (for an event in Ω) or $\mathfrak{c}$ (for an event in Φ), and $\mathfrak{b}$ denotes a bijection between the branches of two isomorphic trees.

2. Evidence

2.1. **Definition of an evidential structure.** The set of an agent's possible evidential states will be denoted by $\mathbb{X}$. This set is a primitive of the theory, so it could be interpreted in a wide variety of ways. The most obvious interpretation, corresponding directly to the idea of a revealed preference theory outlined above, is that an element of $\mathbb{X}$ is something concrete and observable, such as a data set. There are other reasonable interpretations, such as that an evidential state is an equivalence class of data sets, where equivalence means being logically equivalent in the context of some body of background knowledge, of which the agent is certain.[7] For example, a chemist might have been taught that flourine is the only element that forms a binary compound with argon (albeit an unstable one), and also that flourine is the only element that has a boiling point between $-190°$ C and $-187°$ C. Then, the datum that a sample, known to be an element, reacted with argon to form a binary compound would be equivalent to the data set that the sample was solid at $-190°$ C, was gaseous at $-187°$ C, and was liquid at intermediate temperatures. The equivalence class of such data sets would be identified with the evidential state that the sample is flourine.

The elements of $\mathbb{X}$ will be called evidential states (or e-states, for short). There is a binary relation, σ, among e-states, that satisfies assumptions (8)–(12) below. That $x \,\sigma\, y$ (x is weakly more specific than y) means that evidence x is consistent with, and possibly more specific than, evidence y. Σ denotes the strict part of σ and is defined, as usual, by

$$x \,\Sigma\, y \iff [x \,\sigma\, y \text{ and not } y \,\sigma\, x] \tag{3}$$

Also,

$$x \overset{\sigma}{=} y \iff [x \,\sigma\, y \text{ and } y \,\sigma\, x] \tag{4}$$

[7] An interpretation in a similar spirit, regarding an agent who is provided only with summary information about a data set rather than having the opportunity to see the data in its entirety, raises a subtle issue that will be explored further in section 3.4.

Condition (10), below, involves the concept of one state being more specific than the other because of the possession of a single extra datum. E-state x is defined to be immediately more specific than z ($x \,\Sigma_1\, z$) by the condition that

$$x \,\Sigma_1\, z \iff [x \,\Sigma\, z \text{ and not } \exists y \; x \,\Sigma\, y \,\Sigma\, z] \tag{5}$$

Define

$$Y(z) = \{x | x \,\Sigma_1\, z\} \tag{6}$$

Lemma 3 will state that, if $x \,\Sigma\, y$, then there is a finite Σ_1-chain of e-states leading from x to y. That result, which is the most important consequence of assumptions (8) and (10) below, corresponds to the intuitive idea that a data set consists of finitely many reports of observations.

Condition (12) involves the concept of two e-states being incompatible with one another. Incompatibility means that there is no e-state that is demonstrably consistent with both of them, where consistency is demonstrated by being weakly more specific according to σ. Formally, that x and y are incompatible ($x \perp y$) is defined by

$$x \perp y \iff \text{not } \exists z \; [z \,\sigma\, x \text{ and } z \,\sigma\, y] \tag{7}$$

Here are the assumptions regarding e-states and the specificity ordering. A structure satisfying these conditions will be called an evidential structure, or an e-structure.

(8) σ is a preorder (that is, reflexive and transitive). There is no infinite sequence, $z_0, z_1, \ldots$, such that $z_k \,\Sigma\, z_{k+1}$ for all k. That is, σ^{-1} is well founded.

(9) r, an element of $\mathbb{X}$, satisfies $\forall x \; [x \neq \mathsf{r} \implies x \Sigma \mathsf{r}]$, and $\mathbb{X} \neq \{\mathsf{r}\}$.[8]

(10) For any $x \,\Sigma\, z$, there exists y such that $x \,\Sigma_1\, y \,\sigma\, z$.

(11) $Y(x)$ is finite.

(12) If not $x \,\sigma\, z$, then there exists an e-state, y, such that $y \,\sigma\, x$ and $y \perp z$.

2.2. **A classic example.** In Bayesian decision theory, repeated tossing of a coin is a classic example of the accretion of evidence. In that example, e-state r is the empty sequence, corresponding to the situation of the agent before the first toss has been observed. Each e-state is a finite sequence of toss-outcome records (each one, an occurrence of either heads or tails), and $x \,\Sigma_1\, y$ iff x is formed from y by appending just one record, specifying either that the coin landed heads or that it landed tails. More generally, that $x \,\sigma\, y$ means that y is an initial subsequence of x.

Example 1. The set of possible single-outcome records, {'Heads', 'Tails'}, can be regarded as a two-element alphabet. To be slightly more general, let Ξ be a finite alphabet, and define $\mathbb{X} = \bigcup_{n \in \mathbb{N}} \Xi^n$, and specify that $x \,\sigma\, y$ iff y is an initial subsequence of x.

[8] r is the root of $(\mathbb{X}, \sigma)$, regarded as an ordered structure.

Thus defined, $(\mathbb{X}, \sigma)$ satisfies assumptions (8)–(12), and it also possesses three additional features. One of these features is that $(\mathbb{X}, \sigma)$ is a tree. By this statement, we mean that (ξ, Σ_1) is a connected, acyclic, directed graph (with edges pointing toward r, the root), where $x \,\Sigma_1\, y$ indicates that there is a directed edge from x to y.[9] Now, the three features that were just mentioned are as follows.

- Because Ξ is finite, $\mathbb{X}$ is countable.
- Since an e-state is a finite sequence, the length of the sequence serves as a measure of rank, and induction on the rank of an e-state can be used to prove facts about $(\mathbb{X}, \sigma)$.
- For every non-null sequence, x, there is a unique y such that $x \,\Sigma_1\, y$. That is, $(\mathbb{X}, \sigma)$ is a tree.

It will be shown in section 2.3 that every e-structure possesses the first of these features, and that every e-structure possesses a rank function that is imperfectly analogous to sequence length. An e-structure need not be a tree. Subsequently a subclass of the e-structures, to be called experimentation trees, will be defined. Experimentation trees closely mirror the e-structures of example 1 in point of all three features.

Savage represents an agent's beliefs, conditional on e-states described according to example 1 (such as finite coin-toss sequences), in terms of a stochastic process $\langle X_i \colon \Omega \to \Xi \rangle_{i \in \mathbb{N}}$, where $(\Omega, \mathcal{B}, P)$ is a probability space with $\Omega = \Xi^{\mathbb{N}+}$ and $\mathcal{B}$ is the Borel σ-algebra on Ω that is generated by the product of discrete topologies on the copies of Ξ. The pre-image of an e-state, $x \in \Xi^n$, is the event, $\bigcap_{k=1}^{n} X_k^{-1}(x_k)$, and Savage represents the agent's beliefs in e-state x by conditioning P on this event. Let $\mathfrak{e}(x) = \bigcap_{k=1}^{n} X_k^{-1}(x_k)$. For each x, $\mathfrak{e}(x)$ is both open and closed in the product topology on Ω. Also $\mathfrak{e}(x) \subseteq \mathfrak{e}(y)$ if $x \,\sigma\, y$, and $\mathfrak{e}(x) \cap \mathfrak{e}(y) = \emptyset$ if $x \perp y$. Theorem 6, to be stated in section 2.4, asserts that an assignment of events to e-states, satisfying these properties, exists for general e-structures.

2.3. **Every e-structure is countable.** Example 4, specified below, will show that an abstract structure, $(\mathbb{X}, \sigma)$, that satisfies conditions (8)–(12) is not necessarily a tree.[10] However, a rank function can be defined in the general case, and induction on rank proves that $\mathbb{X}$ is a countable set.

The first of the next two lemmas is a straightforward consequence of σ being a preorder according to assumption (8).

Lemma 2. *If $x \overset{\sigma}{=} y$, then, for all z, $z \,\Sigma_1\, x \iff z \,\Sigma_1\, y$.*

[9] See, for example, Gallier [2011, chapter 3] for a more detailed, formal statement of this definition. In section 3, a more restricted class of "experimentation trees" will be specified, and the generic notion of tree that has been defined here will be called a "graph-theoretic tree" when it is necessary to distinguish between the two concepts.

[10] In the example, z is immediately more specific than both y and q.

Lemma 3. *If $x \,\Sigma\, y$, then there exist a function $f\colon \mathbb{N} \to \mathbb{X}$ and a number, n, such that*

$$f(0) = x \text{ and } f(n) = y \text{ and } \forall k{<}n\ f(k)\,\Sigma_1\, f(k+1) \tag{13}$$

There is a function $\rho\colon \mathbb{X} \to \mathbb{N}$ such that, for all x,

$$\begin{aligned}\rho(x) = \min\{m | [x = \mathsf{r} \text{ and } m = 0] \text{ or } \exists g\,[g\colon \mathbb{N} \to \mathbb{X} \text{ and}\\ g(0) = x \text{ and } g(m) = \mathsf{r} \text{ and } \forall k{<}m\ g(k)\,\Sigma_1\, g(k+1)]\}\end{aligned} \tag{14}$$

Proof. A function, f, that satisfies condition (13), will be defined via an auxiliary function, h, that is defined recursively on an initial segment of $\mathbb{N}$. Define $h(0) = x$. Suppose $h(k)$ has been defined and that $h(k) \,\Sigma\, y$. Then $h(k+1)$ will be defined to satisfy $h(k)\,\Sigma_1\, h(k+1)\,\sigma\, y$. Specifically, (10) guarantees that h can be extended in this way.

Since (8) specifies that there is no infinite Σ_1-sequence, there is some $n > 0$ for which $h(n) \stackrel{\sigma}{=} y$, at which stage the recursive process terminates. By lemma 2, $h(n-1)\,\Sigma_1\, y$. Define $f(m) = h(m)$ for $m < n$ and $f(m) = y$ for $m \geq n$. These choices of f and n satisfy (13).

Taking $y = \mathsf{r}$ in (13) shows that the set on the right side of condition (14) is non empty, so equation (14) successfully defines ρ. □

Note that $\rho(\mathsf{r}) = 0$. It might superficially seem from (14) that $z \,\Sigma\, y \implies \rho(z) > \rho(y)$, but the following example shows that the opposite ordering can occur. Nevertheless, induction on the rank of an e-state is a sound method of proof.

Example 4. Let $\mathbb{X} = \{\mathsf{r}, q, s, t, u, v, w, x, y, z\}$ and let σ be the reflexive, transitive closure of $\{(w,\mathsf{r}), (t,w), (x,w), (u,x), (y,x), (v,y), (z,y), (q,\mathsf{r}), (s,q), (z,q)\}$. Then $z \,\Sigma\, y$ but $\rho(z) = 2$ and $\rho(y) = 3$.[11]

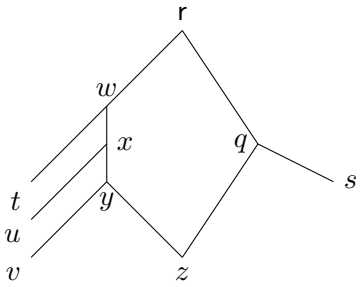

Proposition 5. *For every n, $\{x | \rho(x) \leq n\}$ is finite. $\mathbb{X}$ is countable.*

Proof. Using (9) and (11), it is proved by induction that, for every n, $\{x | \rho(x) \leq n\}$ is finite. Thus the set of all elements $\mathbb{X}$ for which the rank is defined, a countable union of finite sets, is countable. By lemma 3, that set is $\mathbb{X}$ itself. □

[11] E-states t, u, v, and s play no role in the computation of $\rho(y)$ and $\rho(z)$. Rather, their function in the example is to ensure that (12) is satisfied.

2.4. **E-states correspond to events.** Bayesian representation theories identify an e-state with an event, on which probabilities of other, utility-relevant events are conditioned by application of Bayes' theorem. Each event identified with an e-state (and, indeed each event in general), is represented as a subset of a universe of states of nature. Let Ω denote the universe of states of nature. A field of subsets of Ω is a subset, $\mathcal{F}$, of the power set of Ω, such that $\emptyset \in \mathcal{F}$ and $\Omega \in \mathcal{F}$ and $\mathcal{F}$ is closed under union, intersection, and complementation in Ω. For any topology, $\mathcal{T}$, of Ω, define a field of sets, $\mathfrak{F}(\mathcal{T}) = \{B | B \in \mathcal{T} \text{ and } \Omega \setminus B \in \mathcal{T}\}$. For any subset, $\mathcal{S}$, of the power set of Ω, let $\mathcal{T}$ be the topology defined by taking $\mathcal{S}$ to be its sub-base, and define $\mathfrak{F}(\mathcal{S}) = \mathfrak{F}(\mathcal{T})$.[12]

In order to have a minimally satisfactory probability framework, the events identified with e-states must be elements of a field of sets, and the subset relations among events identified with various e-states must mirror the order relations among the e-states themselves. The following theorem, which combines Jech [2002, theorem 14.10] and Stone [1937, theorem 1], asserts that such a nice identification of events with e-states exists.[13]

Theorem 6. *Consider an evidential structure,* $(\mathbb{X}, \sigma)$*. There exist a topological space,* $(\Omega, \mathcal{T})$ *and a mapping,* $\mathfrak{e}\colon \mathbb{X} \to \mathcal{T} \setminus \{\emptyset\}$*, such that, for all* x *and* y *in* $\mathbb{X}$*,*

(15) $(\Omega, \mathcal{T})$ *is compact;*
(16) $\mathfrak{e}(x)$ *is both open and closed;*
(17) $\forall U \in \mathcal{T} \setminus \{\emptyset\} \;\; \exists x \in X \;\; \mathfrak{e}(x) \in U$*;*
(18) $\mathfrak{e}(\mathsf{r}) = \Omega$*;*
(19) $x \,\sigma\, y \iff \mathfrak{e}(x) \subseteq \mathfrak{e}(y)$*;*[14]
(20) $x \perp y \iff \mathfrak{e}(x) \cap \mathfrak{e}(y) = \emptyset$*.*

Moreover,

(21) *every ultrafilter* $\mathcal{U}$ *of* $\mathfrak{F}(\mathcal{T})$ *satisfies* $\mathcal{U} = \{U | \omega \in U \in \mathfrak{F}(\mathcal{T})\}$ *for some point,* ω*.*[15]

[12]Two distinct operators—one on topologies and the other on arbitrary subsets of the power set of Ω—have been defined here. However, there is no ambiguity because, if a set is a topology, then both definitions associate the same field with the set.

[13]Jech's theorem asserts that, under assumptions satisfied by an e-structure, an ordered structure is embeddable in a Boolean algebra, and that there is a minimal Boolean algebra in which it can be embedded. Stone's theorem asserts that a Boolean algebra is isomorphic to $\mathfrak{F}(\mathcal{T})$, where $(\Omega, \mathcal{T})$ is a compact, totally disconnected topological space. A concise exposition of Stone's theorem that makes clear its equivalence to the present theorem 6 is provided by Sikorski [1969, sections I.6–I.8]. Theorem 6 would remain sound, if the finite-branching condition (11) and the well-foundedness assertion in (8) were omitted from the definition of an e-structure. In the present theory, those conditions guarantee that $\mathbb{X}$ is countable. Condition (11) will also be invoked in the proof of lemma 14.

[14]Note that this condition, together with the definitions of Σ and $\overset{\sigma}{=}$, implies that $x \,\Sigma\, y \iff \mathfrak{e}(x) \subset \mathfrak{e}(y)$ and that $x \overset{\sigma}{=} y \iff \mathfrak{e}(x) = \mathfrak{e}(y)$. Also, by (9), $\mathfrak{e}(x) = \Omega \iff x = \mathsf{r}$.

[15]An ultrafilter of $\mathfrak{F}(\mathcal{T})$ is a set of elements of $\mathfrak{F}(\mathcal{T})$ that is closed under pairwise intersection, does not contain $\emptyset$, for every $B \in \mathfrak{F}(\mathcal{T})$, contains either B or $\Omega \setminus B$.

If $(\Phi, \mathcal{S})$ *and* $\mathfrak{c}\colon \Phi \to \mathcal{S} \setminus \{\emptyset\}$*, and also* $(\Omega, \mathcal{T})$ *and* $\mathfrak{e}$*, both satisfy these conditions, then* $\mathfrak{F}(\mathcal{S})$ *and* $\mathfrak{F}(\mathcal{T})$ *are isomorphic as Boolean algebras.*

$(\Omega, \mathcal{T})$ and $\mathfrak{F}(\mathcal{T})$ will be called the canonical sample space and canonical field, respectively, of $(\mathbb{X}, \sigma)$.[16] When $\mathfrak{e}$ is mentioned subsequently in a discussion of a canonical sample space, it should be assumed that $\mathfrak{e}$ satisfies conditions (16)–(20).

2.5. **Embedding an e-structure in a σ-field.** Suppose that $(\Phi, \mathcal{B})$ is a measurable space. That is, suppose that $\mathcal{B}$ is a σ-field on Φ. Let $(\mathbb{X}, \sigma)$ be an evidential structure. A mapping, $\mathfrak{c}\colon \mathbb{X} \to \mathcal{B} \setminus \{\emptyset\}$ is an embedding of $(\mathbb{X}, \sigma)$ into $(\Phi, \mathcal{B})$ if the following conditions are satisfied.[17]

(22) $\quad \mathfrak{c}(\mathsf{r}) = \Phi$ and $\forall x\, \forall y\, [x\, \sigma\, y \iff \mathfrak{c}(x) \subseteq \mathfrak{c}(y)]$

(23) $\quad x \perp y \implies \mathfrak{c}(x) \cap \mathfrak{c}(y) = \emptyset$

(24) $\quad Y(z) \neq \emptyset \implies \bigcup_{x \in Y(z)} \mathfrak{c}(x) = \mathfrak{c}(z)$

A σ-field of particular interest, is the completion of the field generated by the sets of a topology that are both open and closed. That is, define $\mathfrak{C}(\mathcal{T})$ to be the smallest σ-field that contains $\mathfrak{F}(\mathcal{T})$.

Lemma 7. *The mapping,* $\mathfrak{e}$*, posited in theorem 6, embeds* (X, σ) *in* $(\Omega, \mathfrak{C}(\mathcal{T}))$*. If* $\mathfrak{c}(x) = \mathfrak{e}(x) \times \mathbb{X}$*, then* $\mathfrak{c}$ *embeds* (X, σ) *in* $(\Omega \times \mathbb{X}, \mathfrak{C}(\mathcal{T}) \times 2^{\mathbb{X}})$*, where* $\mathfrak{C}(\mathcal{T}) \times 2^{\mathbb{X}}$ *denotes the product* σ*-algebra of* $\mathfrak{C}(\mathcal{T})$ *and the power set of* $\mathbb{X}$*.*

Proof. Conditions (22) and (23) are directly implied by conditions (18), (19), and (20). Since $x\, \sigma\, z$ for each $x \in Y(z)$, (19) implies that $\bigcup_{x \in Y(z)} \mathfrak{e}(x) \subseteq \mathfrak{e}(z)$.

Thus, to establish (24), it is sufficient to show that $\mathfrak{e}(z) \subseteq \bigcup_{x \in Y(z)} \mathfrak{e}(x)$. A contradiction will be derived from the supposition that, to the contrary, $\mathfrak{e}(z) \not\subseteq \bigcup_{x \in Y(z)} \mathfrak{e}(x)$. Then, since $Y(z)$ is finite by (11), and since $\mathfrak{e}(z)$ is open and each $\mathfrak{e}(x)$ is closed, $\mathfrak{e}(z) \setminus \bigcup_{x \in Y(z)} \mathfrak{e}(x)$ is non-empty and open. Therefore, by (17), for some $v \in \mathbb{X}$, $\mathfrak{e}(v) \subseteq \mathfrak{e}(z) \setminus \bigcup_{x \in Y(z)} \mathfrak{e}(x)$.

By (19), $v\, \Sigma\, z$. By lemma 3, there is some w such that $v\, \sigma\, w\, \Sigma_1\, z$. That is, $v\, \sigma\, w$ for some $w \in Y(z)$. By (19), $\mathfrak{e}(v) \subseteq \mathfrak{e}(w) \subseteq \bigcup_{x \in Y(z)} \mathfrak{e}(x)$. But this conclusion contradicts the previous assertion that $\mathfrak{e}(v) \subseteq \mathfrak{e}(z) \setminus \bigcup_{x \in Y(z)} \mathfrak{e}(x)$. Thus $\mathfrak{e}(z) \subseteq \bigcup_{x \in Y(z)} \mathfrak{e}(x)$, so $\mathfrak{e}$ satisfies (24), and hence embeds (X, σ) in $(\Omega, \mathfrak{C}(\mathcal{T}))$.

That $\mathfrak{c}$ (as defined in the lemma) embeds (X, σ) in $(\Omega \times \mathbb{X}, \mathfrak{C}(\mathcal{T}) \times 2^{\mathbb{X}})$ is routinely verified, given that $\mathfrak{e}$ embeds (X, σ) in $(\Omega, \mathfrak{C}(\mathcal{T}))$. □

[16] An element of Ω will be called a 'sample point' as in measure-theoretic probability theory, rather than a 'state of the world' (Savage's terminology), in order to avoid a possible clash of the use of 'state' to denote an e-state—an element of $\mathbb{X}$—and an element of Ω.

[17] Strictly speaking, $\mathfrak{c}$ embeds $(\mathbb{X}, \sigma)$ into the ordered structure $(\mathcal{B}, \subseteq)$. The terminology introduced here is adopted, despite its imprecision, because the measurable space $(\Phi, \mathcal{B})$ is immediately related to the Bayesian rationalization of an agent's pattern of choices.

3. Experimentation trees

3.1. **Definition and example.** The interpretation of x being immediately more specific than y is that x can result from y by the acquisition of a single item of new evidence. The act of acquiring an item of new evidence is an experiment. Subsequently an agent will be envisioned to be acting on the basis of progressive inference, made according to a determinate plan of experimentation. Such a plan is formalized as an experimentation tree.

An experimentation tree in an e-structure, (Ω, σ), consists of a subset $T \subseteq \mathbb{X}$ and a relation $\sigma^* \subseteq \sigma \cap (T \times T)$. Relations Σ^*, Σ_1^*, $\stackrel{*}{=}$, and $\perp^*$ are defined from σ^* by the analogues of equations (3), (5), (4), and (7) respectively. An e-state, $x \in T$, is maximally specific iff

$$\text{not } \exists w \in T \;\; w \,\Sigma^* \, z \tag{25}$$

An experimentation tree in an e-structure, (Ω, σ), is a structure, (T, σ^*), that satisfies:

(26) $\mathsf{r} \in T$ and $T \neq \{\mathsf{r}\}$;
(27) $\sigma^* \subseteq \sigma \cap (T \times T)$ is a preorder of $\mathcal{T}$;
(28) If $x \in T \setminus \{\mathsf{r}\}$, then there is a unique $y \in T$ such that $x \,\Sigma_1^*\, y$;
(29) If $x \,\Sigma_1^*\, y$, then $x \,\Sigma_1\, y$;
(30) If $x \in T$ is not maximally specific, then there are y and z such that $y \,\Sigma_1^*\, x$ and $z \,\Sigma_1^*\, x$ and $y \neq z$.
(31) If $\{x, y, z\} \subseteq T$ and $y \,\Sigma_1^*\, x$ and $z \,\Sigma_1^*\, x$ and $y \neq z$, then $y \perp z$;[18]
(32) If $z \in \mathbb{X}$ and $x \in T$ and $z \,\Sigma\, x$ and x is not a maximally specific node of T, then there exist $v \in \mathbb{X}$ and $w \in T$ such that $v \,\sigma\, z$ and $v \,\sigma\, w$ and $w \,\Sigma_1^*\, x$.

Henceforth, 'tree' will be used to refer to an experimentation tree, and the usual mathematical notion of a tree will be called a 'graph-theoretic tree', where necessary, to avoid ambiguity.[19]

It will be shown below, in corollary 12, that an experimentation tree is is a graph-theoretic tree rooted at r. But the reason why (T, σ^*) is such a tree, is distinct from the reason why $(\mathbb{X}, \sigma)$ is a graph-theoretic tree in the classic example of coin tossing. In that example, there is only one type of experiment to perform. In general, though, there can be several types of experiment that could be performed in various orders. Consider an agent who has received a jar filled with a metalloid compound. The agent can weigh the jar, and can also chemically analyze some of its contents. These distinct experiments cannot be done simultaneously. Therefore, two distinct e-states are immediately less specific than the e-state that the jar contains a kilogram of an arsenic compound: that the jar contains some amount of an arsenic compound, and that the jar contains a kilogram of some

[18] Note that $\perp$, not $\perp^*$, is specified here. The rationale for making this specification will be clear from example 8.

[19] Recall from section 2.5 that $(\mathbb{X}, \sigma)$ is a graph-theoretic tree if $(\mathbb{X}, \Sigma_1)$ is a connected, acyclic, directed graph.

metalloid compound. Under the less-specific-than relation, then, the set of all e-states is not a graph-theoretic tree. However, a completely specified plan of experimentation will dictate the order in which the two experiments will be performed. If the plan specifies that the jar must be weighed before the chemical composition of its contents is determined, then the e-state that the jar contains an unknown quantity of an arsenic compound is not an *immediate* predecessor of the e-state that the jar contains a kilogram of an arsenic compound. That is, the plan of experimentation—represented formally as an experimentation tree—is a graph-theoretic tree because it specifies not only the set of heterogeneous experiments that may be done, but also the specific order in which they would be done.

Condition (27) states that σ^* inherits the reflexivity and transitivity of σ.

Condition (29) formalizes the idea that, if e-state x is reached from e-state y via the conduct of a single experiment, then x must be immediately more specific than y in the underlying evidential structure. That is, essentially, an outcome of a single experiment adds just one datum to the agent's body of evidence.

A node of an experimentation tree, if it is not maximally specific, represents a single experiment that is performed when the agent is in some specific e-state. An experiment must be informative, that is, it cannot have a unique, predetermined outcome. Condition (30) formalizes this idea, that an experiment must lead immediately to at least two distinct, more specific, e-states. Condition (31) formalizes the idea that the possible outcomes of an experiment must be mutually exclusive e-states. Condition (32) can be restated as $[x \,\Sigma\, z \text{ and } z \in T] \implies \neg\forall w \in Y(z)\ x \perp w$. This statement formalizes the idea that, if the agent has reached an e-state at some stage of experimentation, and if another, more specific, e-state cannot be ruled out at that stage, then the agent must not know beforehand that performing the next-stage experiment will necessarily rule out the more specific e-state, regardless of its outcome. Besides providing a concrete example of an experimentation tree, example 8 below shows the rationale for the last two of these conditions.

Example 8. This example concerns e-states regarding the chemical composition of a substance that, in e-state r, is known only to be a compound of one of the metalloid elements germanium (Ge), arsenic (As), or antimony (Sb). Note that arsenic and antimony do, but germanium does not, have a pentoxide compound. The following table lists the e-states in $\mathbb{X}$.

Label	*What the evidence shows the compound to be*
r	Some compound of germanium, arsenic, or antimony.
Ge	Some compound of germanium.
As	Some compound of arsenic.
Sb	Some compound of antimony.
$?O_5$	Arsenic pentoxide or antimony pentoxide.
AsO_5	Arsenic pentoxide.
SbO_5	Antimony pentoxide.
As?	A compound of arsenic, but not arsenic pentoxide.
Sb?	A compound of antimony, but not antimony pentoxide.
??	Either a germanium compound, or a non-pentoxide arsenic or antimony compound.

$(\mathbb{X}, \sigma)$ is shown graphically in first panel of the accompanying figure. Nodes are labelled ends of line segments, the lower end of a segment is immediately more specific than the higher end, and unlabelled intersections of line segments have no significance. In each of the other three panels, a sub-relation of σ determines a graph-theoretic tree on a subset of $\mathbb{X}$. (The deleted elements of σ are shown with dashed lines.) However, only the first two of these are experimentation trees.

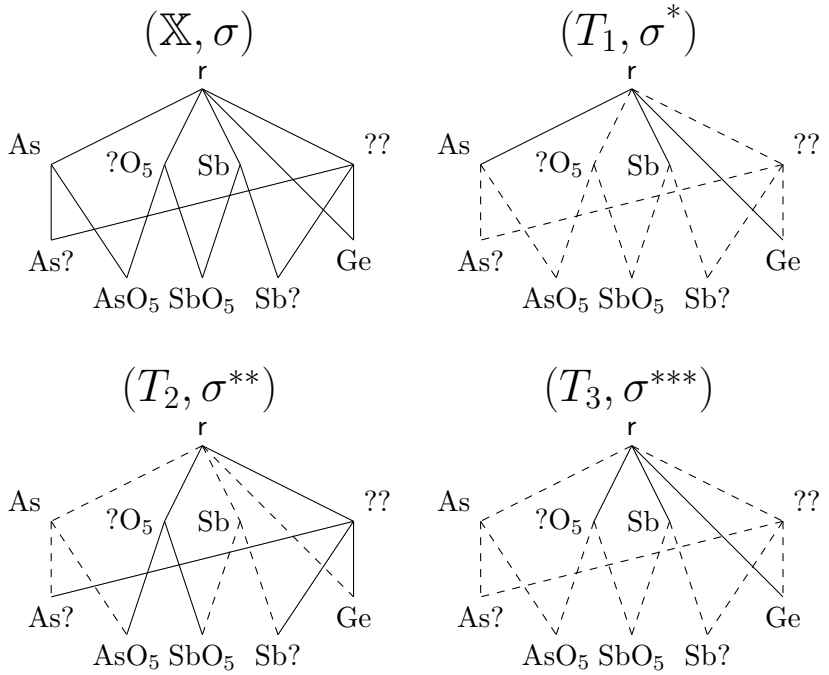

(T_1, σ^*) is an experimentation tree corresponding to a single experiment. The experiment reveals which of the three metalloid elements forms the compound, but does not reveal with what it is combined.

(T_2, σ^{**}) is an experimentation tree corresponding to a two-stage plan of experimentation. At the first stage, it is determined whether or not the

compound is a pentoxide. If so, then, at the second stage, it is determined whether it is arsenic pentoxide or antimony pentoxide. If the compound is not a pentoxide, then, at the second stage, it is determined which metalloid element forms the compound. However, except for the fact that the compound is not a pentoxide, there is no evidence regarding with what the metalloid is combined.

T_3 is formed by deleting As from T_1 and substuting $?O_5$, and σ^{***} is formed by deleting $\mathsf{r}\,\sigma^*$ As from σ^* and substituting $\mathsf{r}\,\sigma^{***}$ $?O_5$. There are two reasons why, despite being a graph-theoretic tree, (T_3, σ^{***}) does not succeed to represent a plan of experimentation. First, it ambigously specifies what the outcome of the experiment would be, if the substance were antimony pentoxide. Second, it does not specify what the result of the experiment would be, if the substance were a non-pentoxide compound of arsenic. These two failures exemplify violations of (31) and (32), respectively. In view of this example, it is clear that (31) and (32) are appropriate conditions to include in the definition of an experimentation tree.

3.2. **Trees are e-structures.** It will be handy to apply results regarding e-structures to experimentation trees. Below, this result will be established as proposition 11.

Lemma 9. *If $x \in T \setminus \{\mathsf{r}\}$, then there exist $m > 0$ and $u_0, \ldots, u_m$ such that $u_0 = x$ and $u_m = \mathsf{r}$ and, for every $i < m$, $u_i \,\Sigma_1^*\, u_{i+1}$.*

Proof. Define $u_0 = x$ and, recursively, if $u_i \neq \mathsf{r}$, then define $u_i \,\Sigma_1^*\, u_{i+1}$. (This defines u_{i+1} uniquely by (28).) By (29), $u_i \,\Sigma_1\, u_{i+1}$. Therefore, by (8), $u_m = \mathsf{r}$ for some m. □

Lemma 10. *If $\{x, y\} \subseteq T$ and not $x \,\sigma^*\, y$ and not $y \,\sigma^*\, x$, then $x \perp y$.*

Proof. $x \neq y$ because σ^* is reflexive. If $x = \mathsf{r}$, then $y \,\Sigma^*\, x$ by lemma 9 and condition (27) (since the asymmetric part of a preorder is transitive), contrary to hypothesis. So $x \neq \mathsf{r}$ and likewise $y \neq \mathsf{r}$.

By lemma 9, there are m and n in $\mathbb{N}$, and sequences $u_0, \ldots, u_m$ and $v_0, \ldots, v_n$, such that $u_0 = x$ and $v_0 = y$ and $u_m = v_n = \mathsf{r}$ and, for $i < m$ and $j < n$, $u_i \,\Sigma_1^*\, u_{i+1}$ and $v_j \,\Sigma_1^*\, v_{j+1}$. Let k be the greatest number such that $u_{m-k} = v_{n-k}$. It must be that $k < \min\{m, n\}$, or else either $x \,\sigma^*\, y$ or $y \,\sigma^*\, x$, contrary to hypothesis.

Let $z = u_{m-k} = v_{n-k}$. Then $u_{m-(k+1)} \,\Sigma_1^*\, z$ and $v_{n-(k+1)} \,\Sigma_1^*\, z$. By (31), $u_{m-(k+1)} \perp v_{n-(k+1)}$. Since $x \,\sigma^*\, u_{m-(k+1)}$ and $y \,\sigma^*\, v_{n-(k+1)}$, $x \perp y$. □

Proposition 11. *If (T, σ^*) is an experimentation tree in $(\mathbb{X}, \sigma)$, then (T, σ^*) is an evidential structure. Its rank function, ρ^*, satisfies*

$$x \,\sigma^*\, y \implies \rho^*(x) \geq \rho^*(y) \tag{33}$$

If $x \,\Sigma_1^*\, y$, then $\rho^*(x) = \rho^*(y) + 1$.

Proof. Let (n*) denote the substitution of σ^* for σ (or of strict relations) in condition (n). It must be proved that σ^* satisfies conditions (8*)–(12*).

Condition (8*) states that σ^* is a preorder (that is, reflexive and transitive); and that there is no infinite sequence, $z_0, z_1, \ldots$, such that $z_k \,\Sigma^*\, z_{k+1}$ for all k (that is, ${\sigma^*}^{-1}$ is well founded). This condition is satisfied, since σ^* is a preorder by (27) and inherits the well foundedness of its inverse from σ.

Condition (9*) states that $\forall x\ [x \neq \mathsf{r} \implies x \Sigma^* \mathsf{r}]$, and that $T \neq \{\mathsf{r}\}$. The property that $\forall x\ [x \neq \mathsf{r} \implies x \,\Sigma\, \mathsf{r}]$ is inherited by σ^* from σ, and (26) asserts that $T \neq \{\mathsf{r}\}$.

Condition (10*) states that, for any $x \,\Sigma^*\, z$, there exists y such that $x \,\Sigma_1^*\, y \,\sigma^*\, z$. To prove it, suppose that $x \,\Sigma^*$z.

If $\mathsf{r} \,\Sigma^*\, z$, then, by (27), $\mathsf{r} \,\sigma\, z$, so not $z \,\Sigma\, \mathsf{r}$. By (9), $z = \mathsf{r}$. But, by (27), this contradicts $\mathsf{r} \,\Sigma^*\, z$, so not $\mathsf{r} \,\Sigma^*\, z$. Since it is supposed that $x \,\Sigma^*\, z$, $x \neq \mathsf{r}$

Therefore, by (28), some $y \in T$ satisfies $x \,\Sigma_1^*\, y$. Now it will be shown by contradiction that $y \,\sigma^*\, z$, proving (10*).

Suppose not $y \,\sigma^*\, z$. Either $z \,\sigma^*\, y$ or not $z \,\sigma^*\, y$. If $z \,\sigma^*\, y$, then $z \,\Sigma^*\, y$, contradicting $x \,\Sigma_1^*\, y$. If not $z \,\sigma^*\, y$, then either $y \neq z$ or $y = z$. If $y \neq z$, then $y \perp z$ by lemma 10, contradicting $x \,\sigma^*\, y$ and $x \,\sigma^*\, z$ (since $\sigma^* \subseteq \sigma$ by (27)). Therefore $y = z$, so $y \,\sigma^*\, z$ by (27).

Condition (11*) states that $\{x | x \,\Sigma_1^*\, y\}$ is finite. This property is inherited from Σ_1, which possesses it by condition (11).

Condition (12*) states that, if $\{x, z\} \subseteq T$ and not $x \,\sigma^*\, z$, then there exists $y \in T$ such that $y \,\sigma^*\, x$ and $y \perp^* z$. If not $x \,\sigma^*\, z$, then either not $z \,\sigma^*\, x$ or $z \,\Sigma^*\, x$. Suppose that not $z \,\sigma^*\, x$. Then $x \perp^* z$ (noting that $\perp \subseteq \perp^*$) by lemma 10, so the condition is satisfied by taking $y = x$. Alternately, suppose that $z \,\Sigma^*\, x$. Define a subset, $\mathcal{S}$, of the power set of T by specifying that $S \in \mathcal{S} \iff [\{x, z\} \subseteq S$ and $\forall u \in S \setminus \{x, z\}\ [z \,\Sigma^*\, u$ and $u \,\Sigma^*\, x]$ and $\forall u \in S\ \forall v \in S\ [u \,\Sigma^*\, v$ or $v \,\Sigma^*\, u]]$. By (8) and (27), there is some w in $\bigcup \mathcal{S}$ such that $w \,\Sigma^*\, x$ and, for all v in $\mathcal{S} \setminus \{w, x\}$, $v \,\Sigma^*\, w$. In particular, either $z = w$ or $z \,\Sigma^*\, w$, so $z \,\sigma^*\, w$.

Now it is shown by contradiction that $w \,\Sigma_1^*\, x$. If, to the contrary, there is some v in T such that $w \,\Sigma^*\, v$ and $v \,\Sigma^*\, x$, then $v \notin \bigcup \mathcal{S}$. But $\bigcup \mathcal{S} \cup \{v\} \in \mathcal{S}$, so $v \in \bigcup \mathcal{S}$. A contradiction has been reached, so $w \,\Sigma_1^*\, x$.

By (30), there is $y \in T$ such that $y \,\Sigma_1^*\, x$ and $w \neq y$. By (31), $y \perp^* w$. For any u, if $u \,\sigma\, w$, then $u \perp y$. In particular, $z \perp^* y$.

Since $z \perp^* y$ if either not $z \,\sigma^*\, x$ or $z \,\Sigma^*\, x$, condition (12*) is satisfied.

At this point, conditions (8*)–(12*) having been verified, it has been established that (T, σ^*) is an evidential structure. Denote its rank function, which lemma 3 guarantees to exist, by ρ^*. Condition (33) will now be proved by contradiction. Suppose that $x \,\sigma^*\, y$ but $\rho^*(x) < \rho^*(y)$. Note that $x \neq y$, since $\rho^*(x) < \rho^*(y)$, so $x \,\sigma^*\, y$ implies that $x \neq \mathsf{r}$.

Let $f\colon \mathbb{N} \to T$ and $g\colon \mathbb{N} \to T$ be functions such as lemma 3 specifies to exist for x and y respectively. That is, $f(0) = x$, $g(0) = y$, $f(\rho^*(x)) =$

$g(\rho^*(y)) = \mathsf{r}$, $f(k)\,\Sigma_1^*\,f(k+1)$ for $k < \rho^*(x)$, and $g(k)\,\Sigma_1^*\,g(k+1)$ for $k < \rho^*(y)$. By (14), there do not exist $h\colon \mathbb{N} \to T$ and $n < \rho^*(y)$ such that $h(0) = y$ and $h(n) = \mathsf{r}$ and $h(k)\,\Sigma_1^*\,h(k+1)$ for $k < n$. In particular, there is no $j \leq \rho^*(x)$ such that $f(j) = y$, or else $h(k) = f(k+j)$ and $n = \rho^*(x) - j$ would be such an impossible path.

Either $y\,\sigma^*\,x$ or not $y\,\sigma^*\,x$.

Suppose that $y\,\sigma^*\,x$. Note that $y \neq \mathsf{r}$. Define $h\colon \mathbb{N} \to T$ by $h(0) = y$ and, for $k > 0$, h(k) = f(k). Define $z = f(1) = h(1)$. Since $\rho^*(y) > \rho^*(x)$, it must be that not $y\,\Sigma_1^*\,z$. But $y\,\Sigma^*\,z$, since $y\,\sigma^*\,x\,\Sigma_1^*\,z$. Therefore, for some w, $y\,\Sigma^*\,w\,\Sigma^*\,z$. But then, since $x\,\sigma^*\,y$, $x\,\Sigma^*\,w\,\Sigma^*\,z$, contradicting $f(0)\,\Sigma_1^*\,f(1)$. Thus, it has been proved by contradiction that not $y\,\sigma^*\,x$.

That is, $x\,\Sigma^*\,y$. By lemma 3, there exist a function $h\colon \mathbb{N} \to T$ and $n \in \mathbb{N}$ such that $h(0) = x$, $h(k+n) = g(k)$ for $0 \leq k \leq \rho^*(y)$, and, for all $k < n + \rho^*(y)$, $h(k)\,\Sigma_1^*\,h(k+1)$.[20] Functions f and h respectively determine two Σ_1^*-paths from x to r, and those paths are distinct because y is on one but not on the other. Since $f(0) = h(0) = y$, there is some $k < \min\{\rho^*(x), n + \rho^*(y)\}$ such that $f(k) = h(k)$ but $f(k+1) \neq h(k+1)$. But this situation contradicts condition (28).

To prove that $x\,\Sigma_1^*\,y \implies \rho^*(x) = \rho(y) + 1$, suppose that $x\,\Sigma_1^*\,y$. Then $\rho^*(x) \leq \rho^*(y) + 1$. Also, by (33), $\rho^*(y) \leq \rho^*(x)$. So, either $\rho^*(x) = \rho(y)$ or $\rho^*(x) = \rho(y) + 1$. By the definition of ρ^*, there is a function, $f\colon \mathbb{N} \to T$, such that $f(0) = x$ and $f(\rho^*(x)) = \mathsf{r}$ and $\forall k < \rho^*(x)\ f(k)\,\Sigma_1^*\,f(k+1)$. It cannot be that $f(1) = y$, or else $\rho^*(y) \leq \rho^*(x) - 1$, contrary to assumption. Therefore, $x\,\Sigma_1^*\,y$ and $x\,\Sigma_1^*\,f(1)$ and $y \neq f(1)$, contrary to condition (28). This proves by contradiction that $\rho^*(x) \neq \rho(y)$, so $\rho^*(x) = \rho(y) + 1$. □

Corollary 12. *If (T, σ^*) is an experimentation tree, then it is a graph-theoretic tree. Relation σ^* is antisymmetric.*

Proof. By lemma 9, (T, Σ_1^*) is a connected graph and every edge is directed toward r. To show that (T, σ^*) is a graph-theoretic tree, it remains to be shown that (T, Σ_1^*) is acyclic. This will be proved by contradiction. Suppose that there were a simple chain in (T, Σ_1^*).[21] There is some e-state, x, of maximum rank among the e-states linked by the edges of the chain. The chain contains edges that link x to two distinct (since the chain is simple) e-states, y and z. By the maximality of $\rho(x)$, and by proposition 11, $x\,\Sigma_1^*\,y$ and $x\,\Sigma_1^*\,z$. This conclusion contradicts (28), so (T, Σ_1^*) must be acyclic.

To show that σ^* is antisymmetric, a contradiction will be derived from the assumptions that $x \neq y$ and $x\,\sigma^*\,y$ and $y\,\sigma^*\,x$. First, from (9) and (27) and $y \neq x$ and $y\,\sigma^*\,x$, it follows that $x \neq \mathsf{r}$. So, by (28), for some z, $x\,\Sigma_1^*\,z$. Since $y\,\sigma^*\,x$, also $y\,\Sigma^*\,z$.

[20] Note that $h(n) = y$ and $h(n + \rho^*(y)) = \mathsf{r}$.

[21] A chain is a simple path in the undirected graph obtained by "forgetting" the directions of edges in (T, Σ_1^*). Cf. Gallier [2011, chapter 3].

Suppose that, for some w, $y\,\Sigma^*\,w$ and $w\,\Sigma^*\,z$. Then, since $x\,\sigma^*\,y$, also $x\,\Sigma^*\,w$, which (together with $w\,\Sigma^*\,z$) contradicts $x\,\Sigma_1^*\,z$. Therefore there can be no such w, so $y\,\Sigma_1^*\,z$.

Since $x\,\Sigma_1^*\,z$ and $y\,\Sigma_1^*\,z$ and $x \neq y$, (31) implies that $x \perp y$. But $x \perp y$ contradicts the assumption that $x\,\sigma^*\,y$. □

Corollary 13. *(T, σ^*) is an experimentation tree in some e-structure, $(\mathbb{X}, \sigma)$, if and only if it is an experimentation tree in itself.*

Proof. The "if" assertion follows immediately from a presumption that (T, σ^*) is an e-structure, which is justified by proposition 11. The "only if" assertion is clear from inspecting the defining conditions of a tree. None of those conditions involves quantification over $\mathbb{X}$, but only over T itself. The only conditions that make reference to the σ relation on $\mathbb{X}$ or to its derived relations are (27), (29), and (31). Those conditions remain true when σ^* is substituted for σ and when $\perp^*$ is substituted for $\perp$. □

In view of corollary 13, an experimentation tree will henceforth be regarded as an e-structure, and denoted by '$(\mathbb{X}, \sigma)$' rather than by '(T, σ^*)', except when (as in section 3.4, below) its status as a substructure of a larger, ambient e-structure is crucial to the discussion.

3.3. **How a tree is related to its canonical field.** The first goal of this section is to relate experimentation trees to partitions and to filtrations, the sequences of progressively finer σ-fields that provide a measure-theoretic setting in which stochastic processes can be defined. Since filtrations are fundamental to the analysis of intertemporal decision making under uncertainty, understanding how a tree induces a filtration provides perspective on research (for example, Epstein and Le Breton [1993], Shmaya and Yariv [2015]) relating "dynamic consistency" to Bayesian rationality.[22]

The second goal is to examine the structure of the canonical field of sets of an experimentation tree. The image under $\mathfrak{e}$ of every e-state in a tree is a block of some finite partition of the canonical sample space into such sets, and the canonical field is the union of the finite fields generated by those partitions. It follows that every element of the canonical field of sets is a finite union of images of e-states. This fact will play a role in the revealed-preference characterization of Bayesian-rational agency.

If $\mathcal{B}$ is a σ-field on Ω, then a filtration of $\mathcal{B}$ is a sequence $\langle \mathcal{B}_n \rangle_{n \in \mathbb{N}}$ of σ-fields such that, for all n, $\mathcal{B}_n \subseteq \mathcal{B}_{n+1} \subseteq \mathcal{B}$.

If $\mathcal{S}$ is a set of subsets of Ω, then let $\mathfrak{C}(\mathcal{S})$ be the smallest σ-field containing $\mathfrak{F}(\mathcal{T})$, where $\mathcal{T}$ is the topology generated by $\mathcal{S}$ as a sub-base.

[22] It is not generally true that an evidential structure induces a filtration that corresponds naturally to its rank function, as will be guaranteed by proposition 15 to happen for experimentation trees. For example, there is clearly no such filtration for the evidential structure of example 4. Implicitly, at least, the literature on dynamically consistent choice under uncertainty models evidence as being an experimentation tree, and some of the conclusions of that literature may not be robust to relaxing that assumption.

A filtration of $\mathfrak{C}(\mathcal{T})$ corresponding to an experimentation tree will be constructed from partitions generated by successive stages of experimentation at nodes of the tree.

Let $\mathfrak{e}$ be the embedding (specified in theorem 6) of an evidential structure into its canonical field of sets. It is first shown that the e-states in an experimentation tree that are immediately more specific than a given e-state in that tree correspond, under $\mathfrak{e}$, to the blocks of a partition of the image of that e-state under $\mathfrak{e}$.

Lemma 14. *If (Ω, σ) is an experimentation tree and z is not maximally specific, then $\{\mathfrak{e}(y)|y \in Y(z)\}$ is a finite partition of $\mathfrak{e}(z)$.*

Proof. Define $\mathcal{Y} = \{\mathfrak{e}(y)|y \in Y(z)\}$. Each $\mathfrak{e}(y)$ is non empty and, by (24) and lemma 7, $\mathfrak{e}(z) = \bigcup \mathcal{Y}$. Thus, to prove that $\mathcal{Y}$ is a partition, it suffices to prove that its elements are disjoint from one another.

Suppose that $x \,\Sigma_1\, z$ and $y \,\Sigma_1\, z$ and $\mathfrak{e}(x) \cap \mathfrak{e}(y) \neq \emptyset$. Then, by (17), there is some w such that $\mathfrak{e}(w) \subseteq \mathfrak{e}(x) \cap \mathfrak{e}(y)$. By (19), $w \,\sigma\, x$ and $w \,\sigma\, y$. That is, not $x \perp y$. It follows from (31) that $x = y$. That is, distinct elements of $\mathcal{Y}$ are disjoint from one another. □

If $(\mathbb{X}, \sigma)$ is an experimentation tree and $\mathfrak{e}\colon \mathbb{X} \to \mathcal{T}$ is according to theorem 6, then define Π_n, for $n \in \mathbb{N}$, by

$$\Pi_n = \{\mathfrak{e}(x)|\rho(x) = n \text{ or } [\rho(x) < n \text{ and } x \text{ is maximally specific}]\} \tag{34}$$

Proposition 15. *If $(\mathbb{X}, \sigma)$ is an experimentation tree and $(\Omega, \mathcal{T})$ is its canonical sample space, then*

$$\mathfrak{e}(\mathbb{X}) = \bigcup_{n\in\mathbb{N}} \Pi_n \tag{35}$$

Each Π_n is a finite partition of Ω, Π_{n+1} refines Π_n, and $\langle \mathfrak{C}(\Pi_n)\rangle_{n\in\mathbb{N}}$ is a filtration of the Borel σ-field of $(\Omega, \mathcal{T})$.[23] *Moreover,*

$$\bigcup_{n\in\mathbb{N}} \mathfrak{C}(\Pi_n) = \mathfrak{F}(\mathcal{T}) \tag{36}$$

Proof. $\mathfrak{e}(\mathbb{X}) = \bigcup_{n\in\mathbb{N}} \Pi_n$ by (34), since every e-state has a finite rank. That each Π_n is a finite partition of Ω into Borel sets is proved by induction. The basis case is that $\Pi_0 = \{\Omega\}$.

Suppose that Π_n is a finite partition of Ω. Let $B = \{x|\rho(x) \leq n$ and x is maximally specific$\}$ and let $C = \{x|\rho(x) = n$ and x is not maximally specific$\}$. Note that $\Pi_n = \{\mathfrak{e}(x)|x \in B \cup C\}$ and that, by proposition 11, $\forall x\ [\rho(x) = n + 1 \implies \exists y\ [y \in C \text{ and } x \,\Sigma_1\, y]]$. Then $\Pi_{n+1} = \{\mathfrak{e}(x)|x \in B\} \cup \{\mathfrak{e}(x)|\rho(x) = n + 1\} = \{\mathfrak{e}(x)|x \in B\} \cup \bigcup_{y\in C}\{\mathfrak{e}(z)|z \,\Sigma_1\, y\}$. By the induction hypothesis and lemma 14, this is a finite partition of Ω that refines Π_n.

[23] Since Π_n is finite, $\mathfrak{C}(\Pi_n) = \mathfrak{F}(\Pi_n)$.

Since Π_n is finite, $\mathfrak{F}(\Pi_n)$ is finite, and is thus (trivially) a σ-field, so $\mathfrak{C}(\Pi_n) = \mathfrak{F}(\Pi_n)$. Since Π_{n+1} refines Π_n, $\mathfrak{C}(\Pi_n) \subseteq \mathfrak{C}(\Pi_{n+1})$. Since $\mathfrak{e}(x) \in \mathcal{T}$ for every x, $\langle \mathfrak{C}(\Pi_n)\rangle_{n\in\mathbb{N}}$ is a filtration of the Borel σ-field of $(\Omega, \mathcal{T})$.

To prove (36), begin by noting that, for all n, $\mathfrak{C}(\Pi_n) \subseteq \mathfrak{F}(\mathcal{T})$, and therefore $\bigcup_{n\in\mathbb{N}} \mathfrak{C}(\Pi_n) \subseteq \mathfrak{F}(\mathcal{T})$. But $\mathfrak{e}(x) \in \bigcup_{n\in\mathbb{N}} \mathfrak{C}(\Pi_n)$ for every x, and $\bigcup_{n\in\mathbb{N}} \mathfrak{C}(\Pi_n) \subseteq \mathfrak{F}(\mathcal{T})$ is a field of sets (because the union of a sequence of successively finer fields is a field), and $\mathfrak{F}(\mathcal{T})$ is the smallest field of sets containing all $\mathfrak{e}(x)$, so $\bigcup_{n\in\mathbb{N}} \mathfrak{C}(\Pi_n) = \mathfrak{F}(\mathcal{T})$. □

Corollary 16. *If $(\Omega, \mathcal{T})$ is the canonical sample space of an experimentation tree, $(\mathbb{X}, \sigma)$, then $\mathfrak{F}(\mathcal{T}) = \{\bigcup_{x\in S} \mathfrak{e}(x) | S \subseteq \mathbb{X}$ and S is finite and $\forall x \in S\ \forall y \in S\ \mathfrak{e}(x) \cap \mathfrak{e}(y) = \emptyset\}$.*

Proof. By proposition 15, $\bigcup_{n\in\mathbb{N}} \mathfrak{C}(\Pi_n) = \mathfrak{F}(\mathcal{T})$. Since Π_n is a finite partition, each block of which is $\mathfrak{e}(x)$ for some $x \in \mathbb{X}$, $\mathfrak{C}(\Pi_n) = \{\bigcup \mathcal{D} | \mathcal{D} \subseteq \Pi_n\} = \{\bigcup_{x\in S} \mathfrak{e}(x) | S \subseteq \mathfrak{e}^{-1}(\Pi_n)\} \subseteq \{\bigcup_{x\in S} \mathfrak{e}(x) | S \subseteq \mathbb{X}$ and S is finite and $\forall x \in S\ \forall y \in S\ \mathfrak{e}(x) \cap \mathfrak{e}(y) = \emptyset\}$. Thus $\mathfrak{F}(\mathcal{T}) \subseteq \{\bigcup_{x\in S} \mathfrak{e}(x) | S \subseteq \mathbb{X}$ and S is finite and $\forall x \in S\ \forall y \in S\ \mathfrak{e}(x) \cap \mathfrak{e}(y) = \emptyset\}$. Since $\mathfrak{e}(x) \in \mathfrak{F}(\mathcal{T})$ for every x, $\{\bigcup_{x\in S} \mathfrak{e}(x) | S \subseteq \mathbb{X}$ and S is finite and $\forall x \in S\ \forall y \in S\ \mathfrak{e}(x) \cap \mathfrak{e}(y) = \emptyset\} \subseteq \mathfrak{F}(\mathcal{T})$. □

3.4. **An e-structure may contain no experimentation tree.** In section 2.1, it was suggested that an evidential state might be taken to be a class of data sets. One such class would be an equivalence class, where equivalence means being logically equivalent in the context of some body of background knowledge that the agent possesses. Another such interpretation would be that an e-state is a class of data sets that are singled out by some partial or summary description of the data. An example will now be given of an evidential structure specified in accordance with this idea, in which there is no experimentation tree.

Example 17. A coin is tossed twice. Consider an agent who does not have the opportunity to observe the two tosses of the coin directly, but who is given information regarding a minimum number of tosses with each outcome. Let $[m, n]$ denote the set of coin-toss sequences in which the coin lands heads at least m times and lands tails at least n times among the two tosses. Then, the set of the agent's possible e-states is $\{[0,0], [0,1], [1,0], [0,2], [1,1], [2,0]\}$, and $\mathsf{r} = [0,0]$. The weakly-more-specific-than relation is defined by $x\,\sigma\,y$ iff every toss-outcome sequence that is consistent with x, is also consistent with y. Equivalently, $[h,i]\,\sigma\,[j,k]$ iff $h \geq k$ and $i \geq k$.

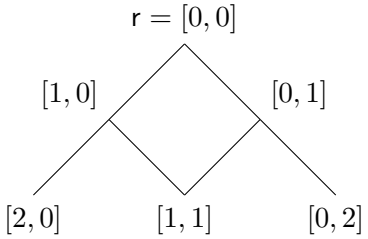

Proposition 18. *There is no experimentation tree in the e-structure of example 17.*

Proof. If (T, σ^*) were an experimentation tree in that e-structure, then, by (26), $\{\mathsf{r}, [m,n]\} \subseteq T$ for some $[m,n] \neq \mathsf{r}$. By (9), $[m,n] \,\Sigma\, \mathsf{r}$, so r is not maximally specific in T. Taking $z = [0,2]$ and $x = \mathsf{r}$ in (32), that condition requires that $[0,1] \,\Sigma_1^*\, \mathsf{r}$. Taking $z = [2,0]$ and $x = \mathsf{r}$, (32) requires that $[1,0] \,\Sigma_1^*\, \mathsf{r}$. But $[1,1] \,\sigma\, [0,1]$ and $[1,1] \,\sigma\, [1,0]$, so not $[0,1] \perp [1,0]$, contradicting (31). □

Various people might draw opposite lessons from proposition 18. Some, who are inclined to believe that some types of evidence may not be discoverable by experimentation, might regard the proposition as formalizing their intuition. Others, who regard the experimental method as a universal engine for the discovery of evidence, might feel that something is wrong here, either with the definition of an e-structure or with the definition of an experimentation tree. Between these two definitions, the authors regard that of an e-structure as being the less self-evident one. One alternative proposal would be to require an e-structure to be a tree. Were that road to be followed, the classic example 1 would seem less special than it does relative to the definition proposed here. But, a liability of that proposal is that example 8 (of an evidential structure regarding metalloid compounds) does not seem peculiar, does contain two experimentation trees, and is not itself a tree. If one would prefer the definition of an e-structure to be tighter than the one presented here is, but would prefer not to require that an e-structure must necessarily be a tree, then the defining conditions (26)–(32) might be regarded as a partially specified definition of an e-structure. As a matter of logic, any proposition of the form, "Every evidential tree in an e-structure satisfies…" that is sound under definition (26)–(32), would remain sound under any more restrictive definition of an e-structure. In particular, theorem 24, the theorem regarding subjective-expected-utility rationalization of a contingent plan specified on an experimentation tree that will be proved below, would remain sound. On the other hand, Examples 20 and 22, which will be formulated later to show that theorem 24 is sound for experimentation trees but not for e-structures in general, might no longer be germane if the definition of an e-structure were to be tightened in some way that would rule them out.

4. Choice alternatives and plans

Plans, to be introduced in this section, are the analogue of demand correspondences in the revealed-preference theory of the consumer. They are the entities, regarding which conditions will be sought that ensure rationalizability in terms of subjective conditional expected utility. A plan specifies which alternative is chosen in each evidential state in some set of evidential states. The set of choice alternatives is a primitive entity in the specification of the theory.

These alternatives should not be confused with acts in Savage's theory. For Savage, the sets of states of the world and of consequences are primitive entities, and an act is constructed as a mapping from states of the world to consequences. The sound analogy between a plan and a Savage act has to do with the aims of the respective theories. The present theory concerns conditions under which a plan can be rationalized in terms of maximization of subjective conditional expected utility, while Savage's theory has to do with conditions under which ordinal preferences among Savage acts can be represented in terms of maximization of subjective expected utility.[24]

Formally, A is the set of choice alternatives.

(37) A is finite and contains at least two distinct alternatives.

If $\emptyset \subset Z \subseteq \mathbb{X}$ and $\zeta\colon Z \to A$, then ζ is a Z-plan. If ζ is an Z-plan for some Z, then ζ is a plan. A tree plan is a plan having some experimentation tree as its domain. If $\emptyset \subset H \subset Z$, ζ is a Z-plan, and η is the restriction of ζ to domain H, then it is said that η is induced by ζ.

5. Conditional preferences, immediate-strict-dominance consistency, and subjective-conditional-expected-utility representation

A conditional preference relation for $\mathbb{X}$ and A is a ternary relation, $R \subseteq A \times A \times \mathbb{X}$, such that

(38) $\preceq_x$ is a weak order (that is, a total preorder) on A, where $a \preceq_x b \iff (a, b, x) \in R$.[25] ($\prec_x$ is the asymmetric part of $\preceq_x$.)

In a discussion where $\mathbb{X}$ and A are clear from context, R will simply be called a conditional preference relation.

A conditional preference relation is immediate-strict-dominance (ISD) consistent iff, for every $z \in \mathbb{X}$,

$$[Y(z) \neq \emptyset \text{ and } \forall x \in Y(z)\ a \prec_x b] \implies a \prec_z b \tag{39}$$

This condition expresses the idea that if, in e-state x, an agent would strictly prefer b to a after acquiring just one piece of further evidence, regardless of what that evidence would show, then the agent should already strictly prefer b to a.

Let R be a conditional preference relation, for A and $\mathbb{X}$. Then R is subjective-conditional-expected-utility (SCEU) representable with respect to $(\mathbb{X}, \sigma)$ if there are an embedding, $\mathfrak{c}$, of $(\mathbb{X}, \sigma)$ into a measurable space, $(\Phi, \mathcal{B})$, and a countably additive probability measure, $\pi\colon \mathcal{B} \to [0, 1]$ and bounded measurable functions, $f_a\colon \Phi \to \mathbb{R}$, for $a \in A$, such that, for each a,

[24] Subjective conditional expected utility has also been analyzed in a Savage-act framework. Cf. Krantz et al. [1971, chapter 8].

[25] In all subsequent discussions, where reference is made to R and $\preceq_x$, this relation between them is assumed.

b, and x,

$$a \preceq_x b \iff \int_{\mathfrak{c}(x)} f_a \, d\pi \leq \int_{\mathfrak{c}(x)} f_b \, d\pi \tag{40}$$

By (40), if $\mathfrak{c}$, π and $\langle J_a \rangle_{a \in A}$ represent R, then

$$[\exists a \, \exists b \; a \prec_x b] \implies \pi(\mathfrak{c}(x)) > 0 \tag{41}$$

Note that

$$\int_B f \, d\pi = \pi(B) \mathsf{E}[f|B] \tag{42}$$

From (41) and (42), it is clear that $f_a(\omega)$ is the state-contingent utility of alternative a at sample point ω.

Regarding why boundedness of f_a is required in the definition of SCEU rationality, three comments are in order.

- SCEU representability is used to define SCEU rationality of a plan. Theorem 24 will characterize SCEU rational plans. That result is proved by a construction of bounded functions. Including boundedness in this definition indirectly incorporates the recognition of that fact in the statement of the theorem.
- Boundedness of the utility function in Savage's representation is implied by Savage's axioms.[26] Thus, incorporating boundedness in this definition would facilitate comparisons between possible results about SCEU representability and their analogues in Savage's framework.
- While some persons view decision theory as being a purely descriptive enterprise of accounting for agents' behavior, others view it as having a role in guiding the formulation and assessment of theories of cognition. On this latter view, the Bayesian theory is a model of agents whose preferences among alternativess are derived from the interaction of two, distinct, mental entities: beliefs and state-contingent utilities. Beliefs are represented by a probability measure. State-contingent utilities are completely distinct from beliefs, so they should induce a preference among acts in combination with any possible probability measure, not only in combination with the probability measure that the agent happens to hold. That is, state-contingent utilities should be integrable with any probability measure, and therefore (in order to avoid a version of the St. Petersburg paradox), the state-contingent utility of any alternative should be a bounded function.

6. SCEU-rational and ISD-consistent plans

Suppose that $(\mathbb{X}, \sigma)$ is an e-structure with canonical sample space $(\Omega, \mathcal{T})$, and that Z is a non-empty subset of $\mathbb{X}$. Conditional preference relation R

[26] Fishburn [1970, theorem 14.5]

for $\mathbb{X}$ and A rationalizes plan $\zeta\colon Z \to A$ with respect to $(\mathbb{X}, \sigma)$ iff, for all $x \in Z$, $\forall a \in A \setminus \{\zeta(x)\}$ $a \prec_x \zeta(x)$.

Plan ζ is SCEU rational with respect to $(\mathbb{X}, \sigma)$ if it is rationalized by a SCEU-representable conditional preference relation for $(\mathbb{X}, \sigma)$ and A.[27] That is, ζ is SCEU-rational iff there exist an embedding, $\mathfrak{c}$, of $(\mathbb{X}, \sigma)$ into a measurable space, $(\Phi, \mathcal{B})$, a probability measure, $\pi\colon \mathcal{B} \to [0,1]$, and, for each $a \in A$, a bounded, measurable function, $f_a\colon \Phi \to \mathbb{R}$, such that, for all $x \in Z$,

$$\forall a \in A \setminus \{\zeta(x)\} \quad \int_{\mathfrak{c}(x)} f_a \, d\pi < \int_{\mathfrak{c}(x)} f_{\zeta(x)} \, d\pi \tag{43}$$

Note that, by (37) and (40) and (43),

$$\text{if } \mathfrak{c},\ \pi \text{ and} \langle f_a \rangle_{a \in A} \text{ SCEU rationalize } \zeta, \text{ then } \forall x\ \pi(\mathfrak{c}(x)) > 0 \tag{44}$$

Call plan $\zeta\colon Z \to A$ ISD consistent if, for every $a \in A$ and $z \in Z$,[28]

$$[Y(z) \neq \emptyset \text{ and } \forall x \in Y(z)\ [x \in Z \text{ and } \zeta(x) = a]] \implies \zeta(z) = a \tag{45}$$

Proposition 19. *A plan is ISD consistent if, and only if, it is rationalized by an ISD-consistent conditional preference relation.*

Proof. Let $\zeta\colon Z \to A$ be a plan in $(\mathbb{X}, \sigma)$. Clearly, if any ISD-consistent conditional preference relation rationalizes ζ, then ζ is an ISD-consistent plan. To prove the converse, define a conditional preference relation, R, by $a \prec_x b \iff [x \in Z \text{ and } a \neq \zeta(x) = b]$. (Otherwise, since $\preceq_x$ is a weak order, both $a \preceq_x b$ and $b \preceq_x a$.) R rationalizes ζ, and ζ is ISD consistent if, and only if, R is ISD consistent. □

A more subtle question is: what is the relation between a plan being SCEU rational and being ISD consistent? It will now be shown that, for plans in general, being ISD consistent is neither necessary nor sufficient for being SCEU rational. It will subsequently be shown that the two conditions are equivalent for tree plans.

The figures below represent plans. A label, '$x \mapsto a$', denotes that node x, at that location, is mapped to alternative a by the plan. In examples 20 and 22, the entire e-structure is the domain of the plan.

Example 20. ISD consistency is not necessary for SCEU rationality. Let $\mathbb{X} = \{\mathsf{r}, x_1, x_2, x_3, z_1, \ldots, z_5\}$. Let Σ_1 comprise the pairs (z_i, x_i), (z_4, x_i), (z_5, x_i), and (x_i, r), and let σ be the reflexive, transitive closure of Σ_1. Let $A = \{a, b\}$. Define plan ζ by $\zeta(\mathsf{r}) = \zeta(z_1) = \zeta(z_2) = \zeta(z_3) = a$ and $\zeta(x_1) = \zeta(x_2) = \zeta(x_3) = \zeta(z_4) = \zeta(z_5) = b$.

[27] This definition corresponds closely to the definition of a Bayes contingent plan in Green and Park [1996], except that here a plan is required to be a function rather than a possibly multi-valued correspondence.

[28] This definition corresponds to the definition of consistency in Green and Park [1996], when ζ (which could be a multi-valued correspondence in that article) is a function.

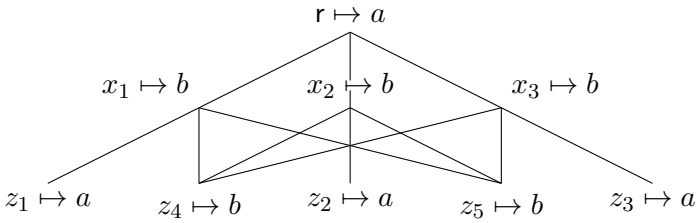

Proposition 21. *Plan ζ of example 20 is SCEU rational but is not ISD consistent.*

Proof. Plan ζ is ISD-inconsistent at r.

By lemma 7, $\mathfrak{e}$ embeds $(\mathbb{X}, \sigma)$ into $(\Omega, \mathfrak{C}(\mathcal{T}))$, where $(\Omega, \mathcal{T})$ is $\{1,2,3,4,5\}$ with the discrete topology, hence $\mathfrak{C}(\mathcal{T})$ is the power set of $\{1,2,3,4,5\}$. Events are assigned by $\mathfrak{e}$ to e-states as follows: $\mathfrak{e}(z_i) = \{i\}$, $\mathfrak{e}(x_i) = \{i,4,5\}$, and $\mathfrak{e}(\mathsf{r}) = \{1,2,3,4,5\}$.

The following probability measure and utility functions constitute an SCEU representation of a coinditional preference relation that rationalizes ζ. π is normalized counting measure on $\{1,2,3,4,5\}$. Utility functions are specified by $f_a(1) = f_a(2) = f_a(3) = f_b(4) = f_b(5) = 1$ and, otherwise, $f_c(\omega) = 0$. □

Example 22. ISD consistency is not sufficient for SCEU rationality. Let $\mathbb{X} = \{\mathsf{r}, x_1, x_2, z_1, z_2, z_3\}$. Let Σ_1 comprise the pairs (z_1, x_1), (z_2, x_2), (z_3, x_1), (z_3, x_2), (x_1, r), and (x_2, r), and let σ be the reflexive, transitive closure of Σ_1. Let $A = \{a, b, c\}$. Define plan ζ by $\zeta(\mathsf{r}) = \zeta(z_3) = a$, $\zeta(x_1) = \zeta(z_2) = b$, and $\zeta(x_2) = \zeta(z_1) = c$.

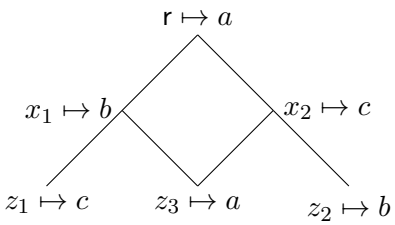

Proposition 23. *Plan ζ of example 22 is ISD consistent but is not SCEU rational.*

Proof. Plan ζ is ISD consistent because there is no e-state, x, for which $\zeta(Y(x))$ is a singleton.

Suppose that ζ were rationalized by a conditional preference relation represented by probability space, $(\Phi, \mathcal{B}, \pi)$, and $\langle f_a \rangle_{a \in A}$, where $\mathfrak{c}$ embeds $(\mathbb{X}, \sigma)$ into $(\Phi, \mathcal{B})$.

The following identities hold, by conditions (23) and (24). $\mathfrak{c}(z_1) \cap \mathfrak{c}(z_2) = \mathfrak{c}(z_1) \cap \mathfrak{c}(z_3) = \mathfrak{c}(z_2) \cap \mathfrak{c}(z_3) = \emptyset$. $\mathfrak{c}(x_i) = \mathfrak{c}(z_i) \cup \mathfrak{c}(z_3)$. $\mathfrak{c}(\mathsf{r}) = \mathfrak{c}(x_1) \cap \mathfrak{c}(x_2)$. It follows that

$$\Phi = \mathfrak{c}(\mathsf{r}) = \mathfrak{c}(x_1) \cup \mathfrak{c}(z_2) = \mathfrak{c}(x_2) \cup \mathfrak{c}(z_1) \tag{46}$$

and

$$\mathfrak{c}(x_1) \cap \mathfrak{c}(z_2) = \mathfrak{c}(x_2) \cap \mathfrak{c}(z_1) = \emptyset \tag{47}$$

Then, since $\zeta(x_1) = \zeta(z_2) = b$ and $\zeta(x_2) = \zeta(z_1) = c$, (43) implies that $\int_{\mathfrak{c}(x_1)} f_b \, d\pi > \int_{\mathfrak{c}(x_1)} f_c \, d\pi$ and $\int_{\mathfrak{c}(z_2)} f_b \, d\pi > \int_{\mathfrak{c}(z_2)} f_c \, d\pi$ and $\int_{\mathfrak{c}(x_2)} f_c \, d\pi > \int_{\mathfrak{c}(x_2)} f_b \, d\pi$ and $\int_{\mathfrak{c}(z_1)} f_c \, d\pi > \int_{\mathfrak{c}(z_1)} f_b \, d\pi$. These inequalities, together with (46) and (47), imply a contradiction: that both $\int_\Phi f_b \, d\pi > \int_\Phi f_c \, d\pi$ and $\int_\Phi f_c \, d\pi > \int_\Phi f_b \, d\pi$. □

7. ISD consistency and SCEU rationality of tree plans

7.1. **The main theorem.** The following theorem will be proved, via several lemmas.[29]

Theorem 24. *Let $(\mathbb{X}, \sigma)$ be an experimentation tree, and consider $\zeta \colon \mathbb{X} \to A$. Plan ζ is SCEU rational with respect to $(\mathbb{X}, \sigma)$ iff is ISD consistent. Moreover, if ζ is SCEU rational, then it is rationalized by a conditional preference relation that is represented by a purely atomic probability measure.*

The assumptions of theorem 24 will be maintained throughout section 7, and it will be assumed that $(\Omega, \mathcal{T})$ is the canonical sample space of $(\mathbb{X}, \sigma)$.

Note that the theorem does not assert that the representing probability measure is unique. For example, there are plans defined on the canonical coin-tossing e-structure of example 1 that it would be natural to rationalize using a nonatomic measure that makes the possible coin-toss sequences to be an exchangeable stochastic process, but that can also be rationalized by purely atomic measures.

7.2. **Branches of $(\mathbb{X}, \sigma)$ and $(\mathfrak{e}(\mathbb{X}), \subseteq)$.** Define a branch of a graph-theoretic tree, $(T, \leq)$, to be a nonempty set, $B \subseteq T$, such that

(48) if $x \in B$ and $x \leq y$, then $y \in B$
(49) if $x \in B$ and $y \in B$, then $x \leq y$ or $y \leq x$
(50) there is no $x \in T \setminus B$ such that, for all $y \in B$, $x < y$

If $(T, \leq)$ is a graph theoretic tree, then let $\mathfrak{B}_T$ be the set of all branches of $(T, \leq)$.

The following result is a consequence of the definitions of an evidential tree and of a branch, of theorem 6 and of corollary 12. Let $\mathbb{E} = \mathfrak{e}(\mathbb{X})$.

[29] A claim that would imply theorem 24 was asserted by Green and Park [1996]. However, their proof was not sound and the claim is false. Specifically, it was supposed that a particular martingale would converge in L_1, but the martingale did not satisfy the sufficient condition (uniform integrability) for such convergence. Eran Shmaya and Leeat Yariv (personal communication) identified the problem and provided a counterexample to the claim.

Lemma 25. *$\mathfrak{e}$ is an isomorphism from $(\mathbb{X}, \sigma)$ to $(\mathbb{E}, \subseteq)$. Thus $(\mathbb{E}, \subseteq)$, is a graph-theoretic tree, and $\mathfrak{e}$ induces a bijection, $\mathfrak{b}\colon \mathfrak{B}_{\mathbb{X}} \to \mathfrak{B}_{\mathbb{E}}$ by*

$$\mathfrak{b}(B) = \{\mathfrak{e}(x) | x \in B\} \tag{51}$$

Lemma 26. *For every $\mathcal{B} \in \mathfrak{B}_{\mathbb{E}}$ and for every n, $\mathcal{B} \cap \Pi_n \neq \emptyset$. Moreover,*

$$\mathcal{U} = \{S | \exists B \in \mathcal{B}\ B \subseteq S \subseteq \Omega\} \tag{52}$$

is an ultrafilter. $\bigcap \mathcal{B}$ is a singleton, $\{\omega\}$. $\mathcal{B} = \{S | \exists n\ \omega \in S \in \Pi_n\}$. For every ω, $\{S | \exists n\ \omega \in S \in \Pi_n\}$ is a branch.

Proof. To prove that $\mathcal{U}$ is an ultrafilter, it must be shown that the intersection of two members of $\mathcal{U}$ also belongs to $\mathcal{U}$, and that for any element S, of $\mathfrak{F}(\mathcal{T})$, either $S \in \mathcal{U}$ or $\Omega \setminus S \in \mathcal{U}$.

If $R \in \mathcal{U}$ and $S \in \mathcal{U}$, then there are $B \in \mathcal{B}$ and $C \in \mathcal{B}$ such that $B \subseteq R$ and $C \subseteq S$. Without loss of generality, assume by (49) that $B \subseteq C$. Then $B \subseteq R \cap S$. Therefore $R \cap S \in \mathcal{U}$.

Suppose that $S \in \mathfrak{F}(\mathcal{T})$. It must be shown that either $S \in \mathcal{U}$ or else $\Omega \setminus S \in \mathcal{U}$.

By proposition 15, there is a minimum n such that $S \in \mathfrak{C}(\Pi_n)$. Let $S = \bigcup \mathcal{S}$, where $\mathcal{S} \subseteq \Pi_n$. Let $\nu = \max\{\rho(B) | B \in \mathcal{B}\} \in \mathbb{N} \cup \{\infty\}$. Either $\nu < n$ or else $n \leq \nu$.

If $\nu < n$, then ν is finite, so, by (49), there is some $B \in \mathcal{B} \cap \Pi_\nu$ such that $\forall C \in \mathcal{B}\ B \subseteq C$. Since $\nu < n$, Π_n refines Π_ν. Let R be a block of Π_n that is a subset of B. R cannot be a strict subset of B, or else (50) would be violated. It must be that B is maximally specific, and $R = B \in \Pi_\nu \cap \Pi_n$. Then, either $B \in \mathcal{S}$ or $B \in \Pi_n \setminus \mathcal{S}$, so either $S \in \mathcal{U}$ or $\Omega \setminus S \in \mathcal{U}$.

If $\nu > n$, then, for some $B \in \mathcal{B}$, $\rho(B) > n$. $B \in \Pi_{\rho(B)}$ by (34) and (35), and $\Pi_{\rho(B)}$ refines Π_n. Therefore, for some $R \in \Pi_n$, either $B \subseteq R \in \mathcal{S}$ or else $B \subseteq R \in \Pi_n \setminus \mathcal{S}$. Correspondingly, either $S \in \mathcal{U}$ or else $\Omega \setminus S \in \mathcal{U}$.

Thus, it has been proved that $\mathcal{U}$ is an ultrafilter and that, for every n, $\mathcal{B} \cap \Pi_n \neq \emptyset$.

By (21) and (52), for some ω, $\bigcap \mathcal{B} = \bigcap \mathcal{U} = \{\omega\}$.

It is routinely verified that $\{S | \exists n\ \omega \in \mathcal{S} \in \Pi_n\}$ is a branch and that, if $\bigcap \mathcal{B} = \{\omega\}$ and $\mathcal{B} \neq \{S | \exists n\ \omega \in S \in \Pi_n\}$, then there is some $B \in \mathcal{B} \setminus \{S | \exists n\ \omega \in S \in \Pi_n\}$. Let $B \in \Pi_n$, and let C be the unique block of Π_n in $\{S | \exists n\ \omega \in \mathcal{S} \in \Pi_n\}$. ($C$ must be unique, by (49).) Then $\omega \notin B$ because $\omega \in C$, contradicting $\bigcap \mathcal{B} = \{\omega\}$. Consequently $\{S | \exists n\ \omega \in S \in \Pi_n\}$ is the unique branch having ω as the intersection of its elements. □

The next result follows immediately from lemma 25 and lemma 26.

Proposition 27. *There is a bijection, $\eta\colon \mathfrak{B}_{\mathbb{X}} \to \Omega$, such that*

$$\{\eta(B)\} = \bigcap \{\mathfrak{e}(z) | z \in B\} \tag{53}$$

Lemma 28. *If ζ is ISD consistent, and if $a \neq \zeta(x)$, then there is a branch, B, of $(\mathbb{X}, \sigma)$ such that*

$$x \in B \text{ and } \forall y \in B\, [y\, \sigma\, x \implies \zeta(y) \neq a] \tag{54}$$

Proof. Define a function, $g\colon \mathbb{N} \to \mathbb{X}$, as follows. For $n \leq \rho(x)$, define $g(n) = f(\rho(x) - n)$, where f is the function defined in lemma 3. Since not $g(n)\, \sigma\, x$ for $n < \rho(x)$, and since $\zeta(x) \neq a$, $g(n)\, \sigma\, x \implies \zeta(g(n)) \neq a$ is satisfied for $n \leq \rho(x)$. Beginning with $n = \rho(x) + 1$, continue defining $g(n)$ recursively in such a way that $g(n)\, \sigma\, x \implies \zeta(g(n)) \neq a$ continues to hold. If $g(n)$ is maximally specific, then set $g(n+1) = g(n)$. Otherwise, since ζ is ISD consistent and $\zeta(g(n)) \neq a$, there is some $y \in Y(x)$ such that $\zeta(y) \neq a$. Set $g(n+1) = y$.

Let $B = g(\mathbb{N})$. By construction, B is a branch of (X, σ) that satisfies (54). □

7.3. **A purely atomic probability space for rationalizing ζ.** Given a consistent plan, ζ, the results of section 7.2 will now be used to construct a probability space, $(\Phi, \mathcal{B}, \pi)$. The probability measure will be concentrated on a countable set of sample points that contains, for every $x \in \mathbb{X}$ and $a \neq \zeta(x)$, a point, ϕ, such that $x \in \eta^{-1}(\phi)$ and $\forall y\, [[y \in \eta^{-1}(\phi) \text{ and } y\, \sigma\, x] \implies \zeta(y) \neq a]$. Subsequently, ζ will be rationalized by constructing state-contingent utility functions, $\langle f_a\colon \Phi \to \mathbb{R} \rangle_{a \in A}$ such that, for each x and $a \neq \zeta(x)$, $f_a(\phi) < f_{\zeta(x)}(\phi)$.

To begin, let $\Phi = \Omega \times \mathbb{X}$ and let $\mathcal{B} = \mathfrak{C}(\mathcal{T}) \times 2^{\mathbb{X}}$. Recall that, by lemma 7, $\mathfrak{c}(x) = \mathfrak{e}(x) \times \mathbb{X}$ defines an embedding of $(\mathbb{X}, \sigma)$ into $(\Omega \times \mathbb{X}, \mathfrak{C}(\mathcal{T}) \times 2^{\mathbb{X}})$. Adopt the convention that, for $\phi \in \Phi$,

$$\phi = (\phi_\Omega, \phi_{\mathbb{X}}) \qquad \phi_\Omega \in \Omega \qquad \phi_{\mathbb{X}} \in \mathbb{X} \tag{55}$$

Define a function, $\theta\colon \mathbb{X} \times A \to \Phi \cup \{\emptyset\}$ by

$$\theta(x,a) = \begin{cases} \emptyset & \text{if } \zeta(x) = a \\ (\eta(B), x) & \text{where } B \text{ satisfies (54), if } \zeta(x) \neq a \end{cases} \tag{56}$$

That is, $\theta(x,a) = (\omega, x)$, where ω corresponds to a branch of $(\mathbb{X}, \sigma)$ that contains x and on which, for any e-state that is weakly more specific than x, plan ζ avoids prescribing alternative a.

Define

$$\Upsilon = \theta(\mathbb{X} \times A) \setminus \{\emptyset\} \tag{57}$$

$\mathbb{X}$ and A are countable sets, so

$$\Upsilon \text{ is a countable subset of } \Phi \tag{58}$$

By (34) and lemma 26, each $\{\omega\}$ is an atom of $\mathfrak{C}(\mathcal{T})$, so each $\{\phi\}$, for $\phi \in \Upsilon$, is an atom of $\mathcal{B}$. Therefore, a function $p\colon \Upsilon \to [0,1]$ such that

$\sum_{\phi\in\Upsilon} p(\phi) = 1$ corresponds to a countably additive, purely atomic, probability measure on $(\Phi, \mathcal{B})$. That is, for every $B \in \mathcal{B}$,

$$\pi(B) = \sum_{\phi\in\Upsilon\cap B} p(\phi) \tag{59}$$

The specific function to be used here is now constructed. By proposition 5 and proposition 11 and (37), there is a bijection, $f\colon \Upsilon \to \mathbb{N}_+$, that satisfies

$$\chi_{\mathbb{X}}\, \Sigma\, \phi_{\mathbb{X}} \implies f(\phi) < f(\chi) \tag{60}$$

Normalizing π to be a probability measure by noting that $\sum_{n=1}^{\infty} 3^{-n} = 1/2$, define

$$p(\phi) = 2 \cdot 3^{-f(\phi)} \tag{61}$$

Since $\sum_{n=k+1}^{\infty} 3^{-n} = 3^{-k}/2$, (60) and (61) imply that, for $\phi \in \Upsilon$,

$$\sum_{\substack{\chi_{\mathbb{X}}\, \Sigma\, \phi_{\mathbb{X}} \\ \chi\in\Upsilon}} p(\chi) < p(\phi) \tag{62}$$

7.4. **Completing the proof of theorem 24.** In order to complete the proof of theorem 24, that an ISD-consistent plan on an experimentation tree is SCEU rational, an embedding, $\mathfrak{c}$, of $(\mathbb{X}, \sigma)$ into $(\Phi, \mathcal{B})$ and bounded, $\mathcal{B}$-measurable functions $\langle f_a\colon \Phi \to \mathbb{R}\rangle_{a\in A}$ must be specified such that (43) is satisfied for all x.

Define

$$\mathfrak{c}(x) = \mathfrak{e}(x) \times \mathbb{X} = \{\phi | \phi_\Omega \in \mathfrak{e}(x)\} \tag{63}$$

and

$$f_b(\phi) = \begin{cases} 1 & \text{if } \phi \in \Upsilon \text{ and } \exists x\, [x\, \sigma\, \phi_{\mathbb{X}} \text{ and } \zeta(x) = b] \\ 0 & \text{otherwise} \end{cases} \tag{64}$$

Consider $\int_{\mathfrak{c}(x)} f_b\, d\pi$. By the definition of π,

$$\int_{\mathfrak{c}(x)} f_b\, d\pi = \sum_{\phi\in\mathfrak{c}(x)\cap\Upsilon} f_b(\phi) p(\phi) \tag{65}$$

Also,

$$\mathfrak{c}(x) \cap \Upsilon = \{\phi | \phi \in \mathfrak{c}(x) \text{ and } \exists y\ \exists a\ [\phi = \theta(y, a)]\} \tag{66}$$

Define subsets, B and C, of $\mathbb{X} \times A$, by

$$\begin{aligned} B &= \{(y, a) | \theta(y, a) \in \mathfrak{c}(x) \text{ and } x\, \sigma\, y\} \\ C &= \{(y, a) | \theta(y, a) \in \mathfrak{c}(x) \text{ and } y\, \Sigma\, x\} \end{aligned} \tag{67}$$

Define

$$\begin{aligned} F &= \{\theta(y, a) | (y, a) \in B\} \\ G &= \{\theta(y, a) | (y, a) \in C\} \end{aligned} \tag{68}$$

Then

$$F \cap G = \emptyset \text{ and } \mathfrak{c}(x) \cap \Upsilon = F \cup G \tag{69}$$

Therefore

$$\int_{\mathfrak{c}(x)} f_b \, d\pi = \sum_{\phi \in F} f_b(\phi) p(\phi) + \sum_{\phi \in G} f_b(\phi) p(\phi) \tag{70}$$

For $\phi \in F$, $y \, \sigma \, x$, so, by (64), $f_{\zeta(x)}(\phi) = 1$.

Let $a \neq \zeta(x)$, and define

$$\begin{aligned} J &= \{\phi | \phi \in F \text{ and } f_a(\phi) = 0\} \\ K &= \{\phi | \phi \in G \text{ and } f_a(\phi) = 1\} \end{aligned} \tag{71}$$

Note that $\theta(x,a) \in J$ and $(\theta(x,a))_{\mathbb{X}} = x$ and $K \subseteq \{\chi | \chi_{\mathbb{X}} \, \Sigma \, x \text{ and } \chi \in \Upsilon\}$. Therefore, by (62),

$$\int_{\mathfrak{c}(x)} f_{\zeta(x)} \, d\pi - \int_{\mathfrak{c}(x)} f_a \, d\pi \;\geq\; p(\theta(x,a)) - \sum_{\substack{\chi_{\mathbb{X}} \, \Sigma \, x \\ \chi \in \Upsilon}} p(\chi) > 0 \tag{72}$$

This inequality shows that condition (43) is satisfied, completing the proof of theorem 24.

References

Bruno de Finetti. Foresight: its logical laws, its subjective sources. *Annales de l'Institut Henri Poincaré*, 7, 1937. translated by Henry E. Kyburg, Jr. in Kyburg and Smokler [1964].

Larry G. Epstein and Michel Le Breton. Dynamically consistent beliefs must be Bayesian. *Journal of Economic Theory*, 61:1–22, 1993.

Peter C. Fishburn. *Utility Theory for Decision Making*. John Wiley & Sons, 1970.

Jean Gallier. *Discrete Mathematics*. Springer, 2011.

Itzhak Gilboa and David Schmeidler. *A Theory of Case-Based Decisions*. Cambridge University Press, 2001.

Edward J. Green and In-Uck Park. Three contributions to the theory of decision under uncertainty. Technical Report 558, Federal Reserve Bank of Minneapolis, 1995. URL `https://www.minneapolisfed.org/research/wp/wp558.pdf`.

Edward J. Green and In-Uck Park. Bayes contingent plans. *Journal of Economic Behavior and Organization*, 31:225–236, 1996. The authors' corrections to proofs were disregarded by the publisher, making the article unreadable. It is advisable to obtain the pre-print in Green and Park [1995].

Thomas J. Jech. *Set Theory*. Springer, third millenium edition, 2002.

Edi Karni. Subjective expected utility without states of the world. *Journal of Mathematical Economics*, 42:325–342, 2006.

Edi Karni. On the definition of states and the choice set in the theory of decision making under uncertainty. 2015. URL `http://krieger.jhu.edu/economics/people/karni/states.pdf`.

Edi Karni, David Schmeidler, and Karl Vind. On state-dependent preferences and subjective probabilities. *Econometrica*, 51:1021–1031, 1983.

David H. Krantz, R. Duncan Luce, Patrick Suppes, and Amos Tversky. *Foundations of Measurement*, volume 1: Additive and Polynomial Representations. Academic Press, 1971.

Henry E. Kyburg, Jr and Howard E. Smokler, editors. *Studies in Subjective Probability*. John Wiley & Sons, 1964.

Matthew Polisson, John K.-H. Quah, and Ludovic Renou. Revealed preferences over risk and uncertainty. Technical Report 740, Oxford University Department of Economics, 2015. URL `http://www.economics.ox.ac.uk/materials/papers/13804/paper740.pdf`.

Marcel K. Richter. Revealed preference theory. *Econometrica*, 34:635–645, 1966.

Leonard Jimmie Savage. The foundations of statistics reconsidered. In *Proceedings of the Fourth Berkeley Symposium on Mathematics and Probability*, volume 1. University of California Press, 1961. reprinted with some changes in Kyburg and Smokler [1964].

Leonard Jimmie Savage. *The Foundations of Statistics*. Dover, second revised edition, 1972.

Eran Shmaya and Leeat Yariv. Experiments on decisions under uncertainty: A theoretical framework. 2015. URL `http://people.hss.caltech.edu/~lyariv/papers/Uncertainty_Experiments.pdf`.

Roman Sikorski. *Boolean Algebras*. Springer-Verlag, third edition, 1969.

M. H. Stone. Applications of the theory of Boolean rings to general topology. *Transactions of the American Mathematical Society*, 41:375–481, 1937.

E-mail address: `fzb5027@psu.edu, eug2@psu.edu`